

A PTAS for Bounded-Capacity Vehicle Routing in Planar Graphs

Amariah Becker^{1(⊠)}, Philip N. Klein¹, and Aaron Schild²

Brown University, Providence, USA
 amariah_becker@brown.edu
 University of California, Berkeley, USA

Abstract. The Capacitated Vehicle Routing problem is to find a minimum-cost set of tours that collectively cover clients in a graph, such that each tour starts and ends at a specified depot and is subject to a capacity bound on the number of clients it can serve. In this paper, we present a polynomial-time approximation scheme (PTAS) for instances in which the input graph is planar and the capacity is bounded. Previously, only a quasipolynomial-time approximation scheme was known for these instances. To obtain this result, we show how to embed planar graphs into bounded-treewidth graphs while preserving, in expectation, the client-to-client distances up to a small additive error proportional to client distances to the depot.

Keywords: Capacitated Vehicle Routing \cdot Approximation algorithms \cdot Metric embeddings

1 Introduction

The Capacitated Vehicle Routing problem with capacity Q>0 for a graph G with client set S and depot r is to find a minimum-cost set of tours that collectively visit every client, such that each tour visits the depot and at most Q clients. This problem arises very naturally in both public and commercial settings including planning school bus routes and package delivery. In general metrics, Capacitated Vehicle Routing is APX-hard, even when Q is a fixed capacity as small as three [1]. In this paper, we show that this hardness result does not extend to planar graphs. Specifically, we give the first polynomial-time approximation scheme (PTAS) for Capacitated Vehicle Routing with fixed capacities in planar graphs.

An embedding of a guest graph G in a host graph H is a mapping ϕ : $V(G) \longrightarrow V(H)$. One seeks embeddings in which, for each pair u, v of vertices of G, the u-to-v distance in G is in some sense approximated by the $\phi(u)$ -to- $\phi(v)$ distance in H. One algorithmic strategy for addressing a metric problem is as

Research supported by NSF grant CCF-1409520. Research supported by NSF Grant CCF-1816861.

[©] Springer Nature Switzerland AG 2019

Z. Friggstad et al. (Eds.): WADS 2019, LNCS 11646, pp. 99-111, 2019.

follows: find an embedding ϕ from the input graph G to a graph H with simple structure; find a good solution in H; lift the solution to a solution in G. The success of this strategy depends on how easy it is to find a good solution in H and how well distances in H approximate corresponding distances in G.

In this paper, we give a randomized method for embedding a planar graph G into a bounded-treewidth host graph H so as to achieve a certain expected distance approximation guarantee. There is a polynomial-time algorithm to find an optimal solution to Bounded-Capacity Vehicle Routing in bounded-treewidth graphs. This algorithm is used to find an optimal solution to the problem induced in H. This solution in the host graph is then lifted to obtain a near-optimal solution in G.

1.1 Related Work

Capacitated Vehicle Routing. There is a substantial body of work on approximation algorithms for Capacitated Vehicle Routing. As the problem generalizes the Traveling Salesman Problem (TSP), for general metrics and values of Q, Capacitated Vehicle Routing is also APX-hard [16]. Haimovich and Rinnoy Kan [14] observe the following lower bound.

$$\frac{2}{Q} \sum_{v \in S} d(v, r) \le cost(OPT) \tag{1}$$

which they use to give a $1 + (1 - \frac{1}{Q})\alpha$ -approximation, where α denotes the approximation ratio of TSP. Using Christofides 1.5-approximation for TSP [9], this gives a $2.5 - \frac{1}{Q}$ approximation ratio. For general metrics and values of Q this result has not been substantially improved upon. Even for tree metrics, the best known approximation ratio for arbitrary values of Q is 4/3, due to Becker [3]. While no polynomial-time approximation schemes are known for arbitrary Q for any nontrivial metric, recently Becker and Paul [7] gave a bicriteria $(1, 1 + \epsilon)$ approximation scheme for tree metrics. It returns a solution of at most the optimal cost, but in which each tour is responsible for at most $(1 + \epsilon)Q$ clients.

One reasonable relaxation is to consider restricted values of Q. Even for Q as small as 3, Capacitated Vehicle Routing is APX-hard in general metrics [1]. On the other hand, for fixed values of Q, the problem can be solved in polynomial time on trees and bounded-treewidth graphs.

Much attention has been given to approximation schemes for Euclidean metrics. In the Euclidean plane \mathbb{R}^2 , PTASs are known for instances in which the value of Q is constant [14], $O(\log n/\log\log n)$ [1], and $\Omega(n)$ [1]. For \mathbb{R}^3 , a PTAS is known for $Q = O(\log n)$ and for higher dimensions \mathbb{R}^d , a PTAS is known for $Q = O(\log^{1/d} n)$ [15]. For arbitrary values of Q, Mathieu and Das designed a quasi-polynomial time approximation scheme (QPTAS) for instances in \mathbb{R}^2 [10]. No PTAS is known for arbitrary values of Q.

Because algorithms for Capacitated Vehicle Routing could be applied to logistics problems in road maps, it is particularly interesting to consider the complexity of approximating the problem in metrics that model road networks.

Becker, Klein, and Saulpic [5] gave a QPTAS for bounded-capacity instances in planar and bounded-genus graphs. The same authors gave a PTAS for graphs of bounded highway dimension [6].

Metric Embeddings. There has been much work on metric embeddings. In particular, Bartal [2] gave a randomized algorithm for selecting an embedding ϕ of the input graph into a tree so that, for any vertices u and v of G, the expected $\phi(u)$ -to- $\phi(v)$ distance in the tree approximates the u-to-v distance in G to within a polylogarithmic factor. Fakcharoenphol, Rao, and Talwar [11] improved the factor to $O(\log n)$.

Talwar [17] gave a randomized algorithm for selecting an embedding of a metric space of bounded doubling dimension and aspect ratio Δ into a graph whose treewidth is bounded by a function that is polylogarithmic in Δ ; the distances are approximated to within a factor of $1+\epsilon$. Feldman, Fung, Könemann, and Post [12] built on this result to obtain a similar embedding theorem for graphs of bounded highway dimension.

What about planar graphs? Chakrabarti et al. [8] showed a result that implies that unit-weight planar graphs cannot be embedded into distributions over $o(\sqrt{n})$ -treewidth graphs so as to achieve approximation to within an $o(\log n)$ factor.

Let us consider distance approximation guarantees with absolute (rather than relative) error. Becker, Klein, and Saulpic [6] gave a deterministic algorithm that, given a constant $\epsilon > 0$, finds an embedding from a graph G of bounded highway dimension to a bounded-treewith graph H such that, for each pair u, v of vertices of G, the $\phi(u)$ -to- $\phi(v)$ distance in H is at least the u-to-v distance in G and exceeds that distance by at most ϵ times the u-to-r distance plus the v-to-r distance, where r is a given vertex of G. This embedding was used to obtain the previously mentioned PTAS for CAPACITATED VEHICLE ROUTING with bounded capacity on graphs of bounded highway dimension.

Recently, Fox-Epstein, Klein, and Schild [13] showed how to embed planar graphs into graphs of bounded treewidth, such that distances are preserved up to a small additive error of ϵD , where D is the diameter of the graph. They show how such an embedding can be used to achieve efficient bicriteria approximation schemes for k-Center and d-Independent Set.

1.2 Main Contributions

In this paper we present the first known PTAS for CAPACITATED VEHICLE ROUTING on planar graphs. We formally state the result as follows.

Theorem 1. For any $\epsilon > 0$ and capacity Q, there is a polynomial-time algorithm for Capacitated Vehicle Routing on planar graphs that returns a solution whose cost is at most $1 + \epsilon$ times optimal.

Prior to this work, only a QPTAS was known [5] for planar graphs. As described in Sect. 1.1, PTASs for CAPACITATED VEHICLE ROUTING are known

only for very few metrics. Our result expands this small list to include planar graphs—a graph class that is quite relevant to vehicle-routing problems as road networks tend to be nearly planar.

The basis for our new PTAS is a new metric-embedding theorem. For a graph G and vertices u and v, let $d_G(u, v)$ denote the u-to-v distance in G.

Theorem 2. There is a constant c and a randomized polynomial-time algorithm that, given a planar graph G with specified root vertex r and given $0 < \epsilon < 1$, computes a graph H with treewidth at most $(\frac{1}{\epsilon})^{c\epsilon^{-1}}$ and an embedding ϕ of G into H, such that, for every pair of vertices u, v of G, $d_G(u, v) \leq d_H(\phi(u), \phi(v))$ with probability 1, and

$$E[d_H(\phi(u), \phi(v))] \le d_G(u, v) + \epsilon[d_G(u, r) + d_G(v, r)] \tag{2}$$

The expectation $E[\cdot]$ is over the random choices of the algorithm.

Why does this metric-embedding result give rise to an approximation scheme for Capacitated Vehicle Routing? We draw on the following observation, which was also used in previous approximation schemes [5,6]: tours with clients far from the depot can accommodate a larger error. In particular, each client can be charged error that is proportional to its distance to the depot. In designing an appropriate embedding, we can afford a larger error allowance for the clients farther from the depot.

Our new embedding result builds on that of Fox-Epstein et al. [13]. The challenge in directly applying their embedding result is that it gives an additive error bound, proportional to the diameter of the graph. This error is too large for those clients close to the depot. Instead, we divide the graph into annuli (bands) defined by distance ranges from the depot and apply the embedding result to each induced subgraph independently, with an increasingly large error tolerance for the annuli farthest from the depot. In this way, each client can afford an error proportional to the diameter of the subgraph it belongs to.

How can these subgraph embeddings be combined into a global embedding with the desired properties? In particular, clients that are close to each other in the input graph may be separated into different annuli. How can we ensure that the embedding approximately preserves these distances while still achieving bounded treewidth?

We draw on a technique that has often been used, e.g. in metric embeddings. We show that by randomizing the choice of where to define the annuli boundaries, and connecting all vertices of all subgraph embeddings to a new, global depot, client distances are approximately preserved (to within their error allowance) in expectation by the overall embedding, without substantially increasing the treewidth. To do so we must ensure that the annuli are wide enough that the probability of nearby clients being separated (and thus generating large error) is small. Simultaneously, the annuli must be narrow enough that, within a given annulus, the clients closest to the depot can afford an error proportional to the distance of the farthest clients from the depot.

A dynamic-programming algorithm can then be used to find an optimal solution to CAPACITATED VEHICLE ROUTING in the bounded-treewidth host graph, and the solution can be lifted to obtain a solution in the input graph that in expectation is near-optimal.

Finally we describe how this result can be derandomized by trying all possible (relevant) choices for defining annuli and noting that for *some* such choice, the resulting solution cost must be near-optimal.

1.3 Outline

In Sect. 2 we describe preliminary notation and definitions. Section 3 describes the details of the embedding and provides an analysis of the desired properties. In Sect. 4 we outline our algorithm and prove Theorem 1. We conclude with some remarks in Sect. 5.

2 Preliminaries

2.1 Basics

Let G = (V, E) denote a graph with vertex set V and edge set E, and let n = |V|. As mentioned earlier, for any two vertices $u, v \in V$, we use $d_G(u, v)$ to denote the length of the shortest u-to-v path in G. We might omit the subscript when the choice of graph is unambiguous. The diameter of a graph G is the maximum distance $d_G(u, v)$ over all choices of u and v.

We say that a graph is *planar* if it can be drawn in the plane without any edge crossings.

We use OPT to denote an optimal solution. For a minimization problem, an α -approximation algorithm is one that returns a solution whose cost is at most α times the cost of OPT. An approximation scheme is a family of $(1+\epsilon)$ -approximation algorithms, indexed by $\epsilon>0$. A polynomial-time approximation scheme (PTAS) is an approximation scheme such that, for each $\epsilon>0$, the corresponding algorithm runs in $O(n^c)$ time, where c is a constant independent of n but may depend on ϵ . A quasi-polynomial-time approximation scheme (QPTAS) is an approximation scheme such that, for each $\epsilon>0$, the corresponding algorithm runs in $O(n^{\log^c n})$ time, where c is a constant independent of n but may depend on ϵ .

An embedding of a guest graph G into a host graph H is a mapping $\phi: V_G \to V_H$ of the vertices of G to the vertices of H.

A tree decomposition of a graph G is a tree T whose nodes (called bags) correspond to subsets of V with the following properties:

- 1. For each $v \in V$, v appears in some bag in T
- 2. For each $(u, v) \in E$, u and v appear together in some bag in T
- 3. For each $v \in V$, the subtree induced by the bags of T containing v is connected

The width of a tree decomposition is the size of the largest bag minus one, and the treewidth of a graph G is the minimum width over all tree decompositions of G.

2.2 Problem Statement

A tour in a graph G is a closed path $v_0, v_1, v_2, ..., v_L$ such that $v_0 = v_L$ and for all $i \in \{1, 2, ..., L\}, (v_{i-1}, v_i)$ is an edge in G.

Given a capacity Q>0 and a graph G=(V,E) with specified client set $S\subseteq V$ and depot vertex $r\in V$, the Capacitated Vehicle Routing problem is to find a set of tours $\Pi=\{\pi_1,\pi_2,...\pi_{|\Pi|}\}$ that collectively cover all clients and such that each tour includes r and covers at most Q clients. The cost of a solution is the sum of the tour lengths, and the objective is to minimize this sum.

If a client s is covered by a tour π , we say that π visits s. Note that π may pass many other vertices (including other clients) that it does not cover.

As stated, the problem assumes that each client has unit demand. In fact, the more general case, where clients have integral demand (assumed to be polynomially bounded) that is allowed to be covered across multiple tours (demand is divisible) reduces to the unit-demand case as follows: For each client $s \in S$ with demand dem(s) = k, add k new vertices $\{v_1, v_2, ..., v_k\}$ each with unit demand and edges (s, v_i) of length zero, and set dem(s) to zero. Note that this modification does not affect planarity. Additionally, since demand is assumed to be polynomially-bounded, the increase in graph size is negligible for the purpose of a PTAS.

For Capacitated Vehicle Routing with indivisible demands, each client's demand must be covered by a single tour, and a tour can cover at most Q units of client demand.

We assume all non-zero distances in G are at least one. If not, the graph can be rescaled. We also assume values of ϵ are less than one. If not, any $\epsilon \geq 1$ can be replaced with a number ϵ' slightly less than one. This only helps the approximation guarantee and does not significantly increase runtime. Of course for very large values of ϵ , an efficient constant-factor approximation can be used instead (see Sect. 1.1).

3 Embedding

In this section, we prove Theorem 2, which we restate for convenience:

Theorem 2. There is a constant c and a randomized polynomial-time algorithm that, given a planar graph G with specified root vertex r and given $0 < \epsilon < 1$, computes a graph H with treewidth at most $(\frac{1}{\epsilon})^{c\epsilon^{-1}}$ and an embedding ϕ of G into H, such that, for every pair of vertices u, v of G, $d_G(u, v) \leq d_H(\phi(u), \phi(v))$ with probability 1, and

$$E[d_H(\phi(u), \phi(v))] \le d_G(u, v) + \epsilon[d_G(u, r) + d_G(v, r)]$$
(3)

The proof uses as a black box the following result from [13]:

Lemma 1 ([13]). There is a number c and a polynomial-time algorithm that, given a planar graph G with specified root vertex r and diameter D, computes a graph H of treewidth at most $(\frac{1}{\epsilon})^c$ and an embedding ϕ of G into H such that, for all vertices u and v,

$$d_G(u, v) \le d_H(\phi(u), \phi(v)) \le d_G(u, v) + \epsilon D$$

For notational convenience, instead of Inequality 3 of Theorem 2, we prove

$$E[d_H(\phi(u),\phi(v))] \le d_G(u,v) + 3\epsilon[d_G(u,r) + d_G(v,r)] \tag{4}$$

from which Theorem 2 can be proved by taking $\epsilon' = \epsilon/3$.

Our embedding partitions vertices of G into bands of vertices defined by distances from r. Choose $x \in [0,1]$ uniformly at random. Let B_0 be the set of vertices v such that $d_G(r,v) < \frac{1}{\epsilon}^{(x)\frac{1}{\epsilon}}$, and for $i \in \{1,2,3,...\}$ let B_i be the set of vertices v such that $\frac{1}{\epsilon}^{(i+x-1)\frac{1}{\epsilon}} \le d_G(r,v) < \frac{1}{\epsilon}^{(i+x)\frac{1}{\epsilon}}$ (see Fig. 1). Let G_i be the subgraph induced by B_i , together with all v-to-v and v-to-v shortest paths for all v-to-v-to-that although the v-to-v-that v-to-v-that diameter of v-that v-to-v-that for v-to-v-that diameter of v-to-v-that is included in v-to-v-that for v-v-that for v-to-v-that for v-that v-to-v-that for v-that v-that for v-that v-

For each G_i , let H_i be the host graph resulting from applying Lemma 1 using $\epsilon' = \epsilon^{\frac{1}{\epsilon}+1}$ and let ϕ_i be the corresponding embedding. Let H be the graph resulting from adding a new vertex r' and for all i and all $v \in B_i$ adding an edge $(\phi_i(v), r')$ of length $d_G(v, r)$. That is, H is formed by connecting (all vertices of) all the H_i to r' (see Fig. 2). Finally, set $\phi(v) = \phi_i(v)$ for all $v \in B_i - \{r\}$ and set $\phi(r) = r'$.

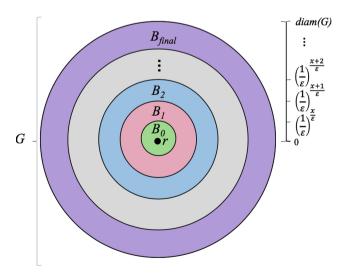


Fig. 1. G is divided into bands $B_0, B_1, ..., B_{final}$ based on distance from r.

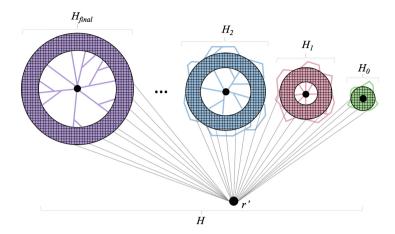


Fig. 2. Each subgraph G_i of G is embedded into a host graph H_i . These graphs are joined via edges to a new depot r' to form a host graph for G.

We can assume that there are at most n bands, since empty bands would not contribute to the embedding. The runtime for constructing H is dominated by the construction of the H_i , which by Lemma [13] is polynomial.

Let H^- be the graph obtained from H by deleting r'. The connected components of H^- are $\{H_i\}_i$. By Lemma 1, the treewidth of each host graph H_i is at most $(\frac{1}{\epsilon'})^{c_0} = (\frac{1}{\epsilon})^{c_0(\epsilon^{-1}+1)}$ for some constant c_0 . This also bounds the treewidth of H^- . Adding a single vertex to a graph increases the treewidth by at most one, so after adding r' back, the treewidth of H is $(\frac{1}{\epsilon})^{c_0(\epsilon^{-1}+1)} + 1 = (\frac{1}{\epsilon})^{c_1\epsilon^{-1}}$ for some constant c_1 .

As for the metric approximation, it is clear that $d_G(u, v) \leq d_H(\phi(u), \phi(v))$ with probability 1. We use the following lemma to prove Inequality 4.

Lemma 2. If $\epsilon d_G(v,r) \leq d_G(u,r) \leq d_G(v,r)$, then the probability that u and v are in different bands is at most ϵ .

Proof. Let i be the nonnegative integer such that $d_G(u,r) = \frac{1}{\epsilon}^{(i+a)\frac{1}{\epsilon}}$ for some $a \in [0,1]$. Let b be the number such that $d_G(v,r) = \frac{1}{\epsilon}^{(i+b)\frac{1}{\epsilon}}$.

$$\frac{1}{\epsilon} \ge \frac{d_G(v,r)}{d_G(u,r)} = \frac{\frac{1}{\epsilon}^{(i+b)\frac{1}{\epsilon}}}{\frac{1}{\epsilon}^{(i+a)\frac{1}{\epsilon}}} = \frac{1}{\epsilon}^{(b-a)\frac{1}{\epsilon}}$$

Therefore

$$b - a \le \epsilon$$

Consider two cases. If $b \le 1$, then the probability that u and v are in different bands is $Pr[a \le x < b] \le \epsilon$.

If b > 1 then the probability that u and v are in different bands is $Pr[x \ge a \text{ or } x \le b-1] \le 1-a+b-1=b-a \le \epsilon$.

We now prove Inequality 4. Let u and v be vertices in G. Without loss of generality, assume $d_G(u,r) \leq d_G(v,r)$. First we address the case where $d_G(u,r) \leq \epsilon d_G(v,r)$. Since $\phi(u)$ and $\phi(v)$ are both adjacent to r' in H, $d_H(\phi(u),\phi(v)) \leq d_H(\phi(u),r') + d_H(\phi(v),r') = d_G(u,r) + d_G(v,r) \leq 2d_G(u,r) + d_G(u,v) \leq d_G(u,v) + 2\epsilon d_G(v,r)$. Therefore $E[d_H(\phi(u),\phi(v))] \leq d_G(u,v) + 3\epsilon[d_G(u,r) + d_G(v,r)]$

Now, suppose $d_G(u,r) > \epsilon d_G(v,r)$. If u and v are in the same band B_i , then by Lemma 1,

$$\begin{split} d_H(\phi(u),\phi(v)) &\leq d_{H_i}(\phi(u),\phi(v)) \leq d_G(u,v) + \epsilon' diam(G_i) \\ &\leq d_G(u,v) + \epsilon' 4\frac{1}{\epsilon}^{(i+x)\frac{1}{\epsilon}} = d_G(u,v) + \epsilon^{\frac{1}{\epsilon}+1} 4\frac{1}{\epsilon}^{(i+x)\frac{1}{\epsilon}} \\ &= d_G(u,v) + \epsilon 4\frac{1}{\epsilon}^{(i+x-1)\frac{1}{\epsilon}} \leq d_G(u,v) + 2\epsilon(d_G(u,r) + d_G(v,r)) \end{split}$$

In the final inequality, when i = 0, we use the fact that all nonzero distances are at least one to give a lower bound on $d_G(u, r)$ and $d_G(v, r)$.

If u and v are in different bands, then since $\phi(u)$ and $\phi(v)$ are both adjacent to r' in H, $d_H(\phi(u), \phi(v)) \leq d_H(\phi(u), r') + d_H(\phi(v), r') = d_G(u, r) + d_G(v, r)$. By Lemma 2, this case occurs with probability at most ϵ .

Therefore $E[d_H(\phi(u),\phi(v))] \leq (d_G(u,v) + 2\epsilon(d_G(u,r) + d_G(v,r))) + \epsilon[d_G(u,r) + d_G(v,r)] \leq d_G(u,v) + 3\epsilon[d_G(u,r) + d_G(v,r)]$, which proves Inequality 4 and completes the proof of Theorem 2.

The construction depends on planarity only via Lemma 1. For the sake of future uses of the construction with other graph classes, we state a lemma.

Lemma 3. Let \mathcal{F} be a family of graphs closed under vertex-induced subgraphs. Suppose that there is a function f and a polynomial-time algorithm that, for any graph G in \mathcal{F} , computes a graph H of treewidth at most $f(\epsilon)$ and an embedding ϕ of G into H such that, for all vertices u and v,

$$d_G(u, v) \le d_H(\phi(u), \phi(v)) \le d_G(u, v) + \epsilon D$$

Then there is a function g and a randomized polynomial-time algorithm that, for any graph G in \mathcal{F} , computes a graph H with treewidth at most $g(\epsilon)$ and an embedding ϕ of G into H, such that, for every pair of vertices u, v of G, with probability $1 d_G(u, v) \leq d_H(\phi(u), \phi(v))$, and

$$E[d_H(\phi(u),\phi(v))] \le d_G(u,v) + \epsilon \left[(d_G(u,r) + d_G(v,r)) \right]$$

4 PTAS for Capacitated Vehicle Routing

In this section, we show how to use the embedding of Sect. 3 to give a PTAS for CAPACITATED VEHICLE ROUTING, proving Theorem 1.

4.1 Randomized Algorithm

We first prove a slight relaxation of Theorem 1 in which the algorithm is randomized, and the solution value is near-optimal *in expectation*. We then show in Sect. 4.2 how to derandomize the result.

Theorem 3. For any $\epsilon > 0$ and capacity Q, there is a randomized algorithm for Capacitated Vehicle Routing on planar graphs that in polynomial time returns a solution whose expected value is at most $1 + \epsilon$ times optimal.

Our result depends on the following lemma, which is proved in the full version [4] of [6].

Lemma 4 (Lemma 20 in [6], **Lemma 15 in** [4]). Given an instance of CAPACITATED VEHICLE ROUTING with capacity Q on a graph G with treewidth w, there is a dynamic-programming algorithm that finds an optimal solution in $n^{O(wQ)}$ time.

Given the dynamic program of Lemma 4 and the embedding of Theorem 2 as black boxes, the algorithm is as follows. First, the graph G is embedded as in Theorem 2 using $\hat{\epsilon} = \epsilon/Q$ into a host graph H with treewidth $(\frac{1}{\hat{\epsilon}})^{c\hat{\epsilon}^{-1}}$ for some constant c, and $d_G(u,v) \leq E[d_H(\phi(u),\phi(v))] \leq d_G(u,v) + \hat{\epsilon}(d_G(u,r) + d_G(v,r))$ for all vertices u and v. The dynamic program of Lemma 4 is then applied to H. The resulting solution SOL_H in H is then mapped back to a solution SOL_G in G which is returned by the algorithm.

Note that the tours in any vehicle-routing solution can be defined by specifying the order in which clients are visited. In particular, we use $(u, v) \in SOL$ to denote that u and v are consecutive elements of {clients} \cup {depot} visited by the solution. In this way, a solution in H is easily mapped back to a corresponding solution in G, as $(u, v) \in SOL_G$ if and only if $(\phi(u), \phi(v)) \in SOL_H$. We use $cost_G(SOL)$ (resp. $cost_H(SOL)$) to denote the cost of a solution SOL in G (resp. H).

We now prove Theorem 3 by analyzing this algorithm.

Lemma 5. For any $\epsilon > 0$ the algorithm described above finds a solution whose expected value is at most $1 + \epsilon$ times optimal.

Proof. Let OPT be the optimal solution in G and let OPT_H be the corresponding induced solution in H. Since the dynamic program finds an optimal solution in H, we have $cost_H(SOL_H) \leq cost_H(OPT_H)$. Additionally, since distances in H are no shorter than distances in G, $cost_G(SOL_G) \leq cost_H(SOL_H)$. Putting these pieces together, we have,

$$E[cost_G(SOL_G)] \leq E[cost_H(SOL_H)] \leq E[cost_H(OPT_H)]$$

$$= E[\sum_{(u,v)\in OPT} d_H(\phi(u),\phi(v))] = \sum_{(u,v)\in OPT} E[d_H(\phi(u),\phi(v))]$$

$$\leq \sum_{(u,v)\in OPT} d_G(u,v) + \hat{\epsilon}(d_G(u,r) + d_G(v,r)) = \sum_{(u,v)\in OPT} d_G(u,v) + 2\hat{\epsilon}\sum_{v\in S} d_G(v,r)$$

$$\leq cost_G(OPT) + 2\hat{\epsilon}\frac{Q}{2}cost_G(OPT) = (1+\epsilon)cost_G(OPT)$$

where the final inequality comes from Lower Bound 1 (see Sect. 1.1).

The following lemma completes the proof of Theorem 3.

Lemma 6. For any $Q, \epsilon > 0$, the algorithm described above runs in polynomial time.

Proof. By Lemma 1, computing H and the embedding of G into H takes polynomial time. By Lemma 4, the dynamic program runs in $|V_H|^{O(wQ)}$ time, where w is the treewidth of H. By Theorem 2, $w = (\frac{1}{\hat{\epsilon}})^{c\hat{\epsilon}^{-1}} = (\frac{Q}{\epsilon})^{c'Q\epsilon^{-1}}$, where c and c' are constants independent of $|V_H|$.

The algorithm therefore runs in $|V_H|^{(Q\epsilon^{-1})^{O(Q\epsilon^{-1})}}$ time. Finally, since $|V_H|$ is polynomial in the size of G, for fixed Q and ϵ , the running time is polynomial.

4.2 Derandomization

The algorithm can be derandomized using a standard technique. The embedding of Theorem 2 partitions the vertices of the input graph into rings depending on a value x chosen uniformly at random from [0,1]. However, the partition depends on the distances of vertices from the root r. It follows that the number of partitions that can arise from different choices of x is at most the number of vertices. The deterministic algorithm tries each of these partitions, finding the corresponding solution, and returns the least costly of these solutions.

In particular, consider the optimum solution *OPT*. As shown in Sect. 4.1,

$$E\left[\sum_{(u,v)\in OPT} d_H(\phi(u),\phi(v))\right]$$

$$= \sum_{(u,v)\in OPT} E\left[d_H(\phi(u),\phi(v))\right]$$

$$< (1+\epsilon)cost_G(OPT).$$

So for some choice of x, the induced cost of OPT in H is nearly optimal, and the dynamic program will find a solution that costs at most as much. This completes the proof of Theorem 1.

5 Conclusion

In this paper, we present the first PTAS for Capacitated Vehicle Routing in planar graphs. Although the approximation scheme takes polynomial time, it is not an *efficient* PTAS (one whose running time is bounded by a polynomial whose degree is independent of the value of ϵ). It is an open question as to whether an efficient PTAS exists. It is also open whether a PTAS exists when the capacity Q is unbounded.

References

- Asano, T., Katoh, N., Tamaki, H., Tokuyama, T.: Covering points in the plane by k-tours: towards a polynomial time approximation scheme for general k. In: Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing, pp. 275–283. ACM (1997)
- Bartal, Y.: Probabilistic approximations of metric spaces and its algorithmic applications. In: 37th Annual Symposium on Foundations of Computer Science, FOCS 1996, Burlington, Vermont, USA, 14–16 October 1996, pp. 184–193 (1996)
- 3. Becker, A.: A tight 4/3 approximation for capacitated vehicle routing in trees. In: Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2018). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2018)
- Becker, A., Klein, P.N., Saulpic, D.: Polynomial-time approximation schemes for k-center and bounded-capacity vehicle routing in metrics with bounded highway dimension. CoRR abs/1707.08270 (2017). http://arxiv.org/abs/1707.08270
- Becker, A., Klein, P.N., Saulpic, D.: A quasi-polynomial-time approximation scheme for vehicle routing on planar and bounded-genus graphs. In: LIPIcs-Leibniz International Proceedings in Informatics, vol. 87. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2017)
- Becker, A., Klein, P.N., Saulpic, D.: Polynomial-time approximation schemes for k-center, k-median, and capacitated vehicle routing in bounded highway dimension. In: 26th Annual European Symposium on Algorithms (ESA 2018). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2018)
- Becker, A., Paul, A.: A PTAS for minimum makespan vehicle routing in trees. arXiv preprint arXiv:1807.04308 (2018)
- 8. Chakrabarti, A., Jaffe, A., Lee, J.R., Vincent, J.: Embeddings of topological graphs: lossy invariants, linearization, and 2-sums. In: 49th Annual IEEE Symposium on Foundations of Computer Science, pp. 761–770 (2008)
- Christofides, N.: Worst-case analysis of a new heuristic for the travelling salesman problem. Technical report, Carnegie-Mellon Univ Pittsburgh Pa Management Sciences Research Group (1976)
- Das, A., Mathieu, C.: A quasi-polynomial time approximation scheme for Euclidean capacitated vehicle routing. In: Proceedings of the Twenty-first Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 390–403. SIAM (2010)
- Fakcharoenphol, J., Rao, S., Talwar, K.: A tight bound on approximating arbitrary metrics by tree metrics. J. Comput. Syst. Sci. 69(3), 485–497 (2004)
- 12. Feldmann, A.E., Fung, W.S., Könemann, J., Post, I.: A $(1+\varepsilon)$ -embedding of low highway dimension graphs into bounded treewidth graphs. In: 42nd International Colloquium on Automata, Languages, and Programming, pp. 469–480 (2015)
- Fox-Epstein, E., Klein, P.N., Schild, A.: Embedding planar graphs into low-treewidth graphs with applications to efficient approximation schemes for metric problems. In: Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1069–1088. SIAM (2019)
- 14. Haimovich, M., Rinnooy Kan, A.: Bounds and heuristics for capacitated routing problems. Math. Oper. Res. 10(4), 527–542 (1985)

- 15. Khachay, M., Dubinin, R.: PTAS for the Euclidean capacitated vehicle routing problem in \mathbb{R}^d . In: Kochetov, Y., Khachay, M., Beresnev, V., Nurminski, E., Pardalos, P. (eds.) DOOR 2016. LNCS, vol. 9869, pp. 193–205. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-44914-2_16
- 16. Papadimitriou, C.H., Yannakakis, M.: The traveling salesman problem with distances one and two. Math. Oper. Res. **18**(1), 1–11 (1993)
- 17. Talwar, K.: Bypassing the embedding: algorithms for low dimensional metrics. In: 36th Annual ACM Symposium on Theory of Computing, pp. 281–290 (2004)